

Rate control protocol for long thin transmission channels

FIELD OF THE INVENTION

The invention deals with a data transmission system comprising at least a transmitter having an adaptable sending rate, a transmission channel having a time varying capacity, and a receiver having data control means for detecting losses and feedback means for reporting said losses to the transmitter, said transmitter having probing means for probing the transmission channel by repeatedly raising its sending rate until a loss is reported by the receiver.

The invention also deals with a receiver to be used in such a transmission system.

The invention also deals with a rate control method to be used for transmitting data from a transmitter having an adaptable sending rate to a receiver designed for detecting losses and reporting losses to said transmitter, through a transmission channel having a time varying capacity, said rate control method comprising the step of repeatedly raising the sending rate of the transmitter for probing the transmission channel until a loss is reported by the receiver.

The invention also deals with a program comprising instructions for implementing the steps of said rate control method that are to be implemented at the receiver, when said program is executed by a processor.

BACKGROUND OF THE INVENTION

Examples of such rate control methods can be found in the article "A survey on TCP-friendly congestion control" by J. Widmer, R. Denda and M. Mauve published in IEEE Network May/June 2001 pages 28 to 37.

For instance, in the TFRC protocol (TCP-Friendly Rate Control protocol), immediately after startup, the transmitter goes into a phase called "slow start" to quickly increase its sending rate to a fair share of the transmission channel bandwidth. Roughly the TFRC slow start consists in doubling the sending rate each RTT (Round Trip Time). The TFRC slow start is terminated with the first loss report.

Such TCP-friendly rate-based control protocols were designed for transmission of non-TCP traffic over the Internet, for instance for the transmission of audio/video streams.

5 An object of the invention is to propose a rate control method specifically adapted to the transmission of such non-TCP traffic over "long thin networks" so as to enable wireless Internet access to audio/video streaming services (long thin networks are defined in the RFC 2757 "Long Thin Networks" by G. Montenegro, S. Dawkins, M. Kojo, V. Magret and N. Vaidya, January 2000).

SUMMARY OF THE INVENTION

10 This is achieved with a data transmission system, a receiver, a rate control method and a program as claimed in claims 1 to 8.

According to the invention, when the transmission channel is being probed by repeatedly raising the sending rate of the transmitter, a fake loss is reported by the receiver when the sending rate has risen to the current capacity of the transmission channel so as to
15 force the transmitter to terminate the probing.

Long thin networks (wireless access links, satellite links, etc...) generally induce large Round Trip Times. Typically the Round Trip Time reaches some hundreds of milliseconds for a wired access to the Internet while it can reach several seconds for a wireless access. This prevents high reactivity. Especially, when the transmission channel is
20 being probed by repeatedly raising the sending rate of the transmitter, there is a risk that the capacity of the transmission channel is exceeded, thereby generating additional losses because the feedback of the receiver is not received on time at the transmitter.

By reporting a fake loss when the sending rate of the transmitter has risen to the current capacity of the transmission channel, the invention overcomes this risk.

25 Advantageously, the receiver computes a received rate that is an estimation of the rate at which data are received, and it monitors the evolution of said received rate for deciding whether the sending rate has risen to the current capacity of the transmission channel.

At the beginning of the probing phase, the received rate increases with the
30 sending rate. It starts to stabilize when the sending rate gets close to the current capacity of the transmission channel. Therefore, monitoring the evolution of the received rate gives an accurate indication of whether the sending rate has risen to the current capacity of the transmission channel.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other aspects of the invention will be further described with reference to the following drawings:

- Figure 1 is a schematic representation of an example of a transmission system according to the invention,
- Figure 2 is a schematic representation of the information that is exchanged in the TFRC protocol,
- Figure 3 is a block diagram of a rate control method according to the invention.

DESCRIPTION OF A PREFERRED EMBODIMENT

A transmission system according to the invention is represented in Figure 1 by way of example. It comprises a transmitter TR, a transmission channel CX, and a receiver RR. In the described embodiment, the transmitter TR is a streaming server. The receiver RR is a mobile apparatus. The transmission channel CX is a long thin transmission channel through a packet transmission network NET (for instance the Internet network) and a radio access network WLS (for instance a network compliant with the GPRS or the UMTS standard). As the transmission channel CX uses shared resources, its capacity can vary in an unpredictable and potentially wide fashion. The transmission system described in Figure 1 allows wireless access to audio/video streaming services provided by the server TR.

According to Figure 1, the transmitter TR comprises a data source DSo, a stream reader SR, a shaper SP, and a transmission / reception block TB. The transmission / reception block TB implements the first four layers of the ISO protocol stack. In this example, the network protocol (ISO layer 3) is IP (Internet Protocol) and the transport protocol (ISO layer 4) is RTP (Real time Transfer Protocol) over UDP (User Datagram Protocol). As UDP is a transport protocol without rate control the transmission / reception block TB comprises a specific rate control block TCB. The aim of the rate control block TCB is to compute a target sending rate $X(t)$ suited to the current capacity of the transmission channel CX. This target sending rate $X(t)$ is sent to the shaper SP. The stream reader SR gets data packets of constant size from the data source DSo. The role of the shaper SP is to control the sending time of these data packets so as to enforce the target sending rate $X(t)$ computed by the rate control block TCB. When delivered by the shaper SP, the data packets are forwarded to the transmission / reception block TB where they are processed according to the above-mentioned transport protocols in order to be transmitted over the transmission channel CX.

The receiver RR comprises a transmission / reception block RB and a stream writer SW. The transmission / reception block RB implements the first four layers of the ISO protocol stack and comprises a rate control block RCB, which functions will be described in more detail in the following of the description. The data packets received over the

5 transmission channel CX are processed by the transmission / reception block RB. Then, they are written in a data sink DSi by the stream writer SW.

In the embodiment described here by way of example, the rate control blocks TCB and RCB implement a rate control method that is compliant with the current version of the TFRC rate control protocol. This is not restrictive: the invention may be used with any
10 other protocol having a probing phase consisting of repeatedly raising the sending rate until a loss is reported by the receiver.

The current version of the TFRC rate control protocol is described in the IETF draft "TCP friendly rate control (TFRC): protocol specification" by M. Handley, J. Padhye S. Floyd and J. Widmer, published on April 27, 2002 and expiring in October 2002.

15 Figure 2 gives a representation of the information that is exchanged between the transmitter TR and the receiver RR according to the TFRC rate control protocol. As described in Figure 2, the receiver RR computes a received rate $R(t)$ and a loss indicator $p(t)$. Both the received rate $R(t)$ and the loss indicator $p(t)$ are transmitted to the transmitter TR in a feedback report (feedback reports are sent from the receiver RR to the transmitter TR by
20 using the RTP control protocol).

The received rate $R(t)$ is an estimation of the rate at which the data have been received since the last feedback report.

The loss indicator $p(t)$ is called loss event rate in TFRC. It is defined as the inverse of the average loss interval, where a loss interval is the number of received packets
25 between two loss events, and a loss event is one or several packets lost within a time interval having the duration of the current Round Trip Time. More details about the computation of the loss event rate can be found in paragraph 5 of the TFRC draft.

When a feedback report is received, the transmitter TR computes a current Round Trip Time $RTT(t)$ and a current Retransmission Time Out $RTO(t)$. The current
30 Retransmission Time Out $RTO(t)$ is derived from the current Round Trip Time $RTT(t)$. For computing the current Retransmission Time Out, TFRC uses the following approximation:
$$RTO(t) = 4 \times RTT(t) .$$

From the current Round Trip Time $RTT(t)$ and the current Retransmission Time Out $RTO(t)$, the transmitter computes an allowed sending rate $T(t)$. The equation to be

used to compute the allowed sending rate is defined in paragraph 3.1 of the TFRC draft and is written as

$$T(t) = \frac{S}{RTT(t) \sqrt{\frac{2p(t)}{3}} + RTO(t) \left[3 \sqrt{\frac{3p(t)}{8}} p(t) (1 + 32p(t)^2) \right]}$$

where S is the mean size of the data packets in bytes.

5 Then the allowed sending rate T(t), the loss indicator p(t) and the received rate R(t) are used to compute the target sending rate X(t). The algorithm to be used for calculating the target sending rate X(t) is defined in paragraph 4.3 of the TFRC draft and is written as :

$$\text{if } p(t) > 0, \text{ then } X(t) = \max \left\{ \min (T(t); 2R(t)); \frac{S}{64} \right\}$$

or if (Now – TimeLastDoubledDuringSlowStart \geq RTT(t)), then

$$10 \quad X(t) = \max \left\{ \min (2X(t-1); 2R(t)); \frac{S}{RTT(t)} \right\}, \text{ and TimeLastDoubledDuringSlowStart} = \text{Now}.$$

“Now” is the time at which the transmitter receives the feedback packet.

Finally the current Round Trip Time RTT(t) is sent from the transmitter TR to the receiver RR because it is used by the receiver for computing the loss indicator p(t).

15 As can be seen from the above, the algorithm used to compute the target sending rate roughly consists in doubling the target sending rate every Round Trip Time when the loss indicator is equal to zero so as to probe the transmission channel CX.

According to the invention, the receiver is designed to report a fake loss (p(t) \neq 0) when the sending rate has risen to the current capacity of the transmission channel CX so as to force the transmitter to terminate said probing. Advantageously with a view to
20 determining if the sending rate has risen to the current capacity of the transmission channel, the receiver advantageously monitors the evolution of the received rate R(t). A preferred embodiment is described in Figure 3.

As indicated in box BX1, the loss indicator p(t) is computed by applying the computation rules indicated in the TFRC draft. Then, in box BX2, a test is performed. If
25 p(t)=0, the execution is pursued as indicated in box BX3. If p(t) \neq 0, a feedback report is directly sent to the transmitter as indicated in box BX6. This feedback report comprises the loss indicator calculated in box BX1.

In box BX3, each time a new value R(t) of the received rate is computed, the receiver RR updates the standard deviation $\sigma(t)$ and the mean value $\mu(t)$ of the last N

values of the received rate $\{R(t), R(t-1), \dots, R(t-N+1)\}$. Then, in box BX4 a test is performed.

If $\sigma(t) < K \cdot \mu(t)$ where K is a constant value, a fake error is generated in box BX5. For instance this is achieved by setting the loss indicator $p(t)$ to a non-zero value. Then
5 a feedback report is sent to the transmitter as indicated in box BX6. This feedback report comprises the loss indicator calculated in box BX5.

If $\sigma(t) \geq K \cdot \mu(t)$, a feedback report is directly sent to the transmitter as indicated in box BX6. This feedback report comprises the loss indicator calculated in box BX1.

10 For instance, $K=0,05$ and $N=3$.

With respect to the above-mentioned transmission system, receiver, rate control method and program, modifications or improvements may be proposed without departing from the scope of the invention. The invention is thus not limited to the examples provided.

15 In particular other or additional criteria may be used by the receiver for deciding if a fake loss has to be generated.

The word "comprising" does not exclude the presence of elements or steps other than those listed in the claims.